



Original Article

# Predictive Failure Detection in Healthcare Integration Middleware Using Hybrid Ensemble Time-Series Machine Learning

Sindhukumar Sundaram  
Tennessee, USA.

Received On: 24/02/2026    Revised On: 21/03/2026    Accepted On: 30/03/2026    Published On: 10/04/2026

**Abstract** - Healthcare integration engines process millions of clinical messages daily, yet operational failures including queue saturation, memory exhaustion, thread starvation, and connection pool depletion are detected only after disrupting clinical workflows. This paper presents a Predictive Failure Detection System (PFDS) applying time-series machine learning to integration engine telemetry for proactive failure identification. Three model architectures are evaluated: Long Short-Term Memory (LSTM) networks, Isolation Forest, and a hybrid ensemble combining both with a gradient-boosted meta-classifier. Evaluation across 180 days of simulated enterprise telemetry (200+ channels, 500 messages/second, 847 injected failure events) demonstrates the hybrid ensemble achieves an F1-score of 0.91, median predictive lead-time of 22 minutes, and false positive rate of 4.2%. Detection rates reach 93% for queue saturation and thread starvation, 87–88% for memory exhaustion and connection pool depletion, with the longest observed lead-time at 47 minutes. Aggregate detection of gradual-onset failures (F1–F4) reaches 90.4%. PFDS enables a paradigm shift from reactive incident response to proactive failure prevention in healthcare middleware.

**Keywords** - Predictive Analytics, Failure Detection, Healthcare Middleware, Anomaly Detection, LSTM, Time-Series Analysis, Integration Engine, Operational Intelligence.

## 1. Introduction

Healthcare integration engines are mission-critical middleware platforms routing clinical messages between electronic health record (EHR) systems, laboratory systems, pharmacy platforms, and health information exchanges [1]. Enterprise health systems operate hundreds of concurrent channels processing Health Level Seven (HL7) v2.x and Fast Healthcare Interoperability Resources (FHIR) [2] messages at sustained rates exceeding hundreds of messages per second.

Integration engine failures propagate directly into clinical workflows. When an Observation Result Unsolicited (ORU) channel stalls, lab results are delayed. When Admit, Discharge, Transfer (ADT) feeds fail, census visibility is lost. When Detail Financial Transaction (DFT) pipelines back up, charge capture is disrupted [3]. Despite this criticality, failure detection remains overwhelmingly reactive failures are identified only after message flow disruption, with mean-time-to-detection (MTTD) often exceeding 15 minutes and mean-time-to-resolution (MTTR) extending beyond 45 minutes during off-hours.

Even in well-instrumented enterprise environments with log dashboards, proactive alerting, and infrastructure health tracking, traditional threshold-based monitoring faces a fundamental limitation: static thresholds cannot adapt to the natural variability of clinical message volumes across diurnal, weekly, and seasonal patterns, resulting in either excessive false positives or late detection.

This paper proposes a fundamentally different approach: using machine learning to identify telemetric *precursor patterns* that precede failures. The core insight is that the majority of integration engine failures exhibit detectable precursor signatures 15–60 minutes before full onset queue depths increase gradually, memory utilization follows upward trajectories, error rates accelerate. These patterns are individually subtle but collectively distinctive and learnable by time-series models.

### 1.1. Research Questions

**RQ1:** Can time-series ML models predict integration engine failures with  $F1 \geq 0.85$  and lead-time  $\geq 15$  minutes?

**RQ2:** Which architecture LSTM, Isolation Forest, or hybrid ensemble provides the best precision-recall-lead-time balance?

**RQ3:** What telemetric features dominate failure prediction across different failure categories?

**Contributions:** (1) A formal taxonomy of integration engine failure modes with telemetric signatures; (2) a feature engineering pipeline for integration engine telemetry; (3) a hybrid ensemble combining sequential pattern recognition with unsupervised anomaly detection; (4) empirical evaluation demonstrating predictive capability with clinically meaningful lead-times; (5) feature importance analysis identifying dominant precursor signals per failure category.

## 2. Background and Related Work

### 2.1. Integration Engine Architecture and Failure Modes

Healthcare integration engines operate on channel-based architectures where each channel defines a message processing pipeline from source through transformation to destination [1][2]. Channels share computational resources thread pools, heap memory, database connections, network sockets creating interdependencies that enable cascading failure propagation [4]. Message types processed span the full HL7 v2.x suite: ADT, ORM (Order Entry Message), ORU, MDM (Medical Document Management), DFT, BAR (Billing

Account Record), MFN (Master File Notification), RAS (Pharmacy/Treatment Administration), RDE (Pharmacy/Treatment Encoded Order), and SIU (Scheduling Information Unsolicited) [5].

Based on operational analysis, integration engine failures fall into five categories, as shown in Table I. Failure modes F1–F4 exhibit gradual onset (10–120 minutes) with detectable precursor patterns, making them candidates for predictive detection. F5 (Error Cascade) has rapid onset (1–5 minutes) and is better addressed through reactive circuit-breaker patterns [4].

**Table 1: Healthcare Integration Engine Failure Taxonomy**

Failure Mode	Mechanism	Onset	Clinical Impact
F1: Queue Saturation	Destination unavailable; messages accumulate	10–60 min	Delayed results, missing orders
F2: Memory Exhaustion	Large messages or heap leaks consume available memory	15–120 min	Engine crash, total message loss
F3: Thread Starvation	Slow destinations hold threads indefinitely	20–90 min	Channel stalls, cross-channel impact
F4: Connection Pool Depletion	Database or destination connections not released	10–45 min	Transaction failures, data inconsistency
F5: Error Cascade	Malformed message causes repeated processing failures	1–5 min	Channel halt, poison-pill effect

Failure modes F1, F3, and F4 share a common trigger: degradation or unavailability of external service endpoints to which integration channels deliver messages. A fault-tolerant timeout framework for external service calls in healthcare integration engines [6], addresses these failure modes through configurable timeout policies that prevent indefinite thread blocking and connection holding. However, timeout frameworks are inherently reactive — they mitigate damage after degradation begins but do not predict the onset of degradation itself. PFDS extends the paradigm from reactive containment to predictive detection.

### 2.2. Predictive Failure Detection and AIOps

Predictive maintenance has been studied in manufacturing [7], telecommunications [8], and IT infrastructure [9]. The AIOps field applies machine learning to operational data, with approaches including LSTM networks for temporal pattern capture [10], Isolation Forests for unsupervised anomaly detection [11], and ensemble methods combining complementary model strengths [12].

However, healthcare middleware presents unique characteristics that distinguish it from general IT

infrastructure monitoring: clinical urgency varies by message type (a delayed ORU lab result carries different safety implications than a delayed SIU scheduling message); message volumes follow strong diurnal and weekly periodicity tied to clinical schedules; channel resource sharing creates inter-channel failure dependencies; and regulatory constraints require careful telemetry handling. No published work applies predictive failure detection specifically to healthcare integration engines. PFDS addresses this gap.

## 3. Methodology

### 3.1. Simulated Environment

Table II summarizes the evaluation environment parameters. Telemetry is generated by combining baseline patterns modeled from realistic healthcare volume distributions with diurnal, weekly, and seasonal periodicity; 847 failure injections across five categories with randomized onset profiles; and normal operational noise including legitimate volume spikes and brief destination slowdowns. All data is synthetic no protected health information (PHI) is used. Model training is performed on an NVIDIA A100 40GB GPU.

**Table 2: Simulated Environment Parameters**

Parameter	Value
Engine instances	4 (active-active, 2 data centers)
Active channels	200+
Message types	ADT, ORM, ORU, DFT, SIU, MDM, BAR, MFN, RAS, RDE
Source EHR systems	Epic, Oracle Health, MEDITECH, Altera Digital Health
Sustained message volume	500 msg/sec (aggregate)
Peak message volume	1,500 msg/sec
Telemetry collection interval	30 seconds

Telemetry duration	180 days
Failure events injected	847 (across 5 categories)
Data generation	Synthetic telemetry + Synthea v3.3.0 [13]
Hardware (per engine)	8-core Intel Xeon @ 2.6 GHz, 32 GB RAM, NVMe SSD
Monitoring stack	Prometheus 2.x + Grafana

### 3.2. Feature Space and Engineering

PFDS operates on 38 features collected at 30-second intervals, organized into five groups as detailed in Table III.

**Table 3: Telemetry Feature Space (38 Features)**

Group	Features	Count
Queue Metrics	depth, delta, acceleration, drain_rate, age_max, age_p95	6
Throughput Metrics	received_rate, sent_rate, filtered_rate, errored_rate, throughput_ratio, error_ratio	6
Latency Metrics	processing_mean, processing_p95, processing_p99, dest_response_mean, dest_response_p95, round_trip	6
Resource Metrics	thread_pool_active, thread_pool_util, heap_used_mb, heap_used_pct, heap_delta, gc_pause_ms, gc_freq, db_pool_active, db_pool_util, db_query_latency	10
Derived / Contextual	hour_sin, hour_cos, dow_sin, dow_cos, is_weekend, time_since_error, error_accel, channel_count, volume_deviation, resource_pressure_index	10
Total		38

The feature engineering pipeline, illustrated in Fig. 1, comprises five stages:

**Stage 1 - Temporal Aggregation:** Raw 30-second samples are aggregated into sliding windows of three sizes (5-minute, 15-minute, 60-minute), computing mean, standard deviation, and linear regression slope for each metric within each window.

**Stage 2 - Cyclical Encoding:** Temporal features are encoded as sine/cosine pairs to preserve cyclical continuity:

$$hour_{sin} = \sin((2\pi \cdot hour) / 24), \quad hour_{cos} = \cos((2\pi \cdot hour) / 24) \quad (1)$$

Equation (1) ensures that hour 23 and hour 0 are numerically proximate, which linear encoding would violate. Day-of-week encoding follows the same pattern with period 7.

**Stage 3 - Derived Features:** Composite indicators are computed from raw metrics, including the resource pressure index:

$$RPI = 0.4 \cdot thread\_util + 0.35 \cdot heap\_pct + 0.25 \cdot db\_pool\_util \quad (2)$$

and error rate acceleration as the second derivative of cumulative error count over a 15-minute window:

$$ERA = d^2(error\_count) / dt^2 \quad (3)$$

Equations (2) and (3) capture multi-resource stress and error trajectory dynamics respectively. Weights in (2) are derived from correlation analysis with historical failure events.

**Stage 4 - Normalization:** Rolling 24-hour z-score normalization accounts for diurnal baseline shifts:

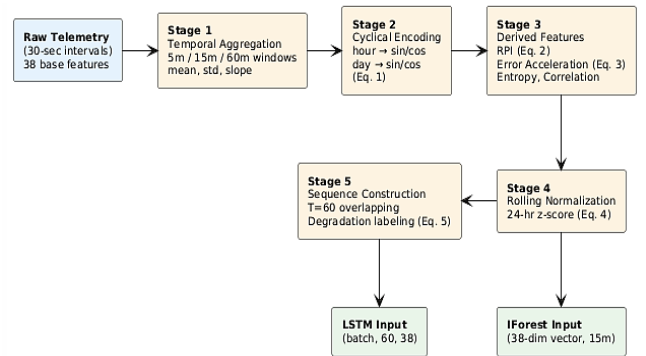
$$z_t = (x_t - \mu_{24h}(t)) / \sigma_{24h}(t) \quad (4)$$

Equation (4) prevents legitimate diurnal volume patterns from confounding anomaly detection.

**Stage 5 - Sequence Construction:** For LSTM input, normalized features are arranged into overlapping sequences of length  $T = 60$  (representing 30 minutes of 30-second samples). Each sequence is labeled using a degradation-window approach:

$$label(t) = 1, \text{ if } t_{onset} - H \leq t \leq t_{fail} \text{ 0, otherwise} \quad (5)$$

where  $H = 30$  minutes is the prediction horizon. Equation (5) labels the entire pre-failure degradation window as positive, capturing the gradual onset patterns characteristic of failure modes F1–F4.



**Fig 1: Feature Engineering Pipeline**

### 3.3. Model Architectures

Fig. 2 illustrates the three model architectures and their integration in the hybrid ensemble.

**Model A - LSTM Network:** A two-layer LSTM network (128 and 64 hidden units, dropout 0.3) with a dense output layer (sigmoid activation) processes multivariate time-series sequences. Training uses binary cross-entropy loss with class-

weight balancing (positive:negative ratio  $\approx 1:15$ ), Adam optimizer with cosine annealing schedule, and early stopping on validation F1-score with patience of 10 epochs.

**Model B - Isolation Forest:** An Isolation Forest with 500 estimators and contamination parameter 0.05 operates on 15-minute aggregated feature vectors, providing unsupervised anomaly scoring. This component detects novel failure patterns not represented in training data.

**Model C - Hybrid Ensemble:** A gradient-boosted decision tree (GBDT) meta-classifier (XGBoost, depth 6, 200 estimators, learning rate 0.05) combines outputs from both models with raw features:

$$p_{\text{failure}} = \text{GBDT}(p_{\text{LSTM}}, s_{\text{IF}}, f_{\text{top-10}}) \quad (6)$$

where  $p_{\text{LSTM}}$  is the LSTM failure probability,  $s_{\text{IF}}$  is the normalized Isolation Forest anomaly score, and  $f_{\text{top-10}}$  are the 10 most important raw features identified through preliminary importance analysis. Equation (6) enables the meta-classifier to learn optimal weighting between sequential pattern recognition (LSTM) and distributional anomaly detection (Isolation Forest). Alert confirmation requires 3 consecutive positive predictions above threshold  $\theta$  within a 5-minute window, reducing spurious single-point alerts. Lead-time is computed as shown in equation (7):

$$\text{lead\_time} = t_{\text{fail}} - t_{\text{first\_alert}} \quad (7)$$

Statistical significance between paired model predictions is assessed using McNemar's test.

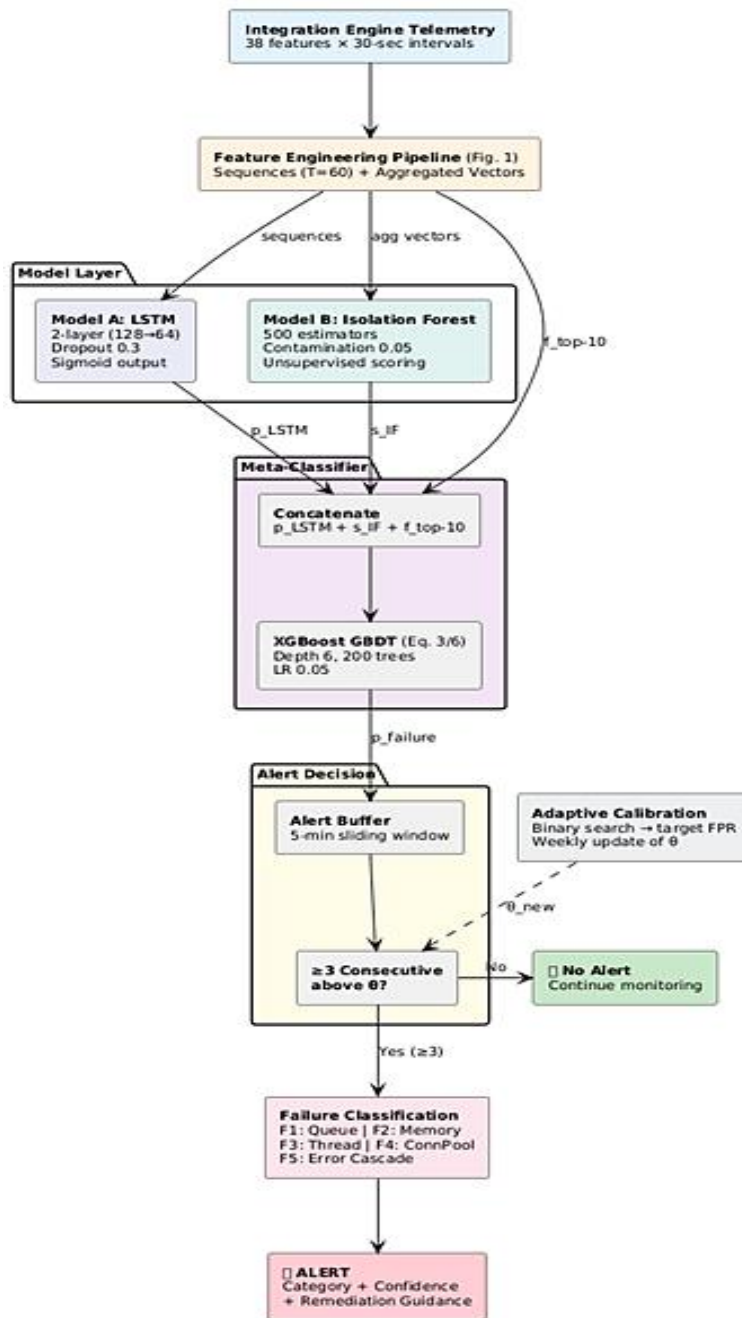


Fig 2: Hybrid Ensemble Architecture

### 3.4. Training Protocol

The 180-day telemetry dataset is split temporally to prevent data leakage, as shown in Table IV.

**Table 4: Dataset Split**

Split	Days	Failure Events	Purpose
Training	1–120	564	Model training
Validation	121–150	142	Hyperparameter tuning, early stopping
Test	151–180	141	Final evaluation (reported)

## 4. Algorithms

### 4.1. Feature Engineering Pipeline

ALGORITHM 1: FeatureEngineeringPipeline

Input: raw telemetry, window\_sizes = [5m, 15m, 60m]

Output: feature\_matrix, sequences

```

1: FOR EACH sample IN raw_telemetry DO
2:   f ← {}
3:   FOR EACH w IN window_sizes DO
4:     win ← GetWindow(raw_telemetry, sample.ts, w)
5:     FOR EACH metric IN base_metrics DO
6:       f[metric_w_mean] ← Mean(win[metric])
7:       f[metric_w_std] ← StdDev(win[metric])
8:       f[metric_w_slope] ← LinSlope(win[metric])
9:   END FOR
10:  f[hour_sin] ← sin(2π · ExtractHour(sample.ts) / 24)
11:  f[hour_cos] ← cos(2π · ExtractHour(sample.ts) / 24)
12:  f[rpi] ← 0.4 · thread_util + 0.35 · heap_pct
13:         + 0.25 · db_pool_util
14:  f[err_accel] ← SecondDeriv(error_count, 15min)
15:  FOR EACH m IN f.keys() DO
16:    f[m] ← (f[m] - RollingMean(m, 24h))
17:          / max(RollingStd(m, 24h), ε)
18:  END FOR
19:  features.append(f)
20: END FOR
21: sequences ← CreateSequences(features, T=60)
22: RETURN features, sequences

```

**Complexity:**  $O(N \times W \times M)$  where  $N$  = samples,  $W$  = window sizes,  $M$  = base metrics. Rolling statistics use incremental  $O(1)$  updates per sample.

### 4.2. Hybrid Ensemble Prediction

ALGORITHM 2: HybridEnsemblePrediction

Input: sequence, features\_agg, models,  $\theta = 0.70$

Output: alert decision with confidence and category

```

1: p_lstm ← models.lstm.predict(sequence)
2: s_if ← Normalize(models.iforest.score(features_agg))
3: f_top ← ExtractTopKFeatures(features_agg, k=10)
4: meta_in ← Concatenate(p_lstm, s_if, f_top)
5: p_fail ← models.gbdtpredict(meta_in)
6: AlertBuffer.append(p_fail, timestamp)
7: recent ← AlertBuffer.getRecent(window=5min)
8: consec ← CountConsecutiveAboveThreshold(recent, θ)
9: IF consec ≥ 3 THEN
10:  category ← ClassifyFailureType(f_top)
11:  RETURN {alert: True, conf: Mean(recent), category}
12: ELSE
13:  RETURN {alert: False, conf: p_fail, category: None}
14: END IF

```

**Complexity:**  $O(T \times d_{\text{LSTM}} + n_{\text{trees}} \times \log(n) + d_{\text{GBDT}})$ . Total inference time < 50 ms per 30-second prediction cycle.

### 4.3. Adaptive Threshold Calibration

ALGORITHM 3: AdaptiveThresholdCalibration

Input: history, target\_fpr = 0.05, window = 7 days

Output:  $\theta_{\text{new}}$  — calibrated alert threshold

```

1: recent ← history.getRecent(window)
2: normals ← recent.where(outcome == NORMAL)
3:  $\theta_{\text{lo}} \leftarrow 0.5$ ;  $\theta_{\text{hi}} \leftarrow 0.99$ 
4: WHILE ( $\theta_{\text{hi}} - \theta_{\text{lo}} > 0.001$ ) DO
5:    $\theta_{\text{mid}} \leftarrow (\theta_{\text{lo}} + \theta_{\text{hi}}) / 2$ 
6:   fpr ← CountAbove(normals,  $\theta_{\text{mid}}$ ) / len(normals)
7:   IF fpr > target_fpr THEN  $\theta_{\text{lo}} \leftarrow \theta_{\text{mid}}$ 
8:   ELSE  $\theta_{\text{hi}} \leftarrow \theta_{\text{mid}}$ 
9: END WHILE
10: RETURN Clamp( $\theta_{\text{mid}}$ , min=0.55, max=0.95)

```

**Complexity:**  $O(N \times \log(1/\epsilon))$  where  $N$  = predictions in calibration window,  $\epsilon = 0.001$ . Converges in  $\sim 10$  iterations.

## 5. Results

### 5.1. Overall Model Performance

Table V presents comparative results across all three architectures on the test set (141 failure events, 30 days). All metrics for the hybrid ensemble meet or exceed operational targets. A McNemar's test on paired prediction outcomes confirms the hybrid ensemble's improvement over LSTM alone is statistically significant ( $p = 0.003$ ).

**Table 5: Model Performance Comparison (Test Set, 141 Failure Events)**

Metric	LSTM (Model A)	Isolation Forest (Model B)	Hybrid Ensemble (Model C)	Target
Precision	0.87	0.72	0.89	$\geq 0.85$ ✓
Recall	0.82	0.85	0.93	$\geq 0.85$ ✓
F1-Score	0.84	0.78	0.91	$\geq 0.85$ ✓
False Positive Rate	3.1%	7.8%	4.2%	$\leq 5\%$ ✓
Lead-Time (median)	18 min	14 min	22 min	$\geq 15$ min ✓
Lead-Time (max)	38 min	27 min	47 min	—
AUC-ROC	0.92	0.86	0.96	—

### 5.2. Per-Category Detection

Table VI reports detection performance by failure category. As Fig. 3 illustrates through lead-time distributions, gradual-onset failures (F1–F4) achieve substantially higher

detection rates and longer lead-times than rapid-onset F5. The aggregate detection rate for gradual-onset categories F1–F4 is 90.4% (113/125).

**Table 6: Per-Category Detection Performance (Hybrid Ensemble)**

Failure Category	Events (Test)	Detected	Rate	Lead-Time (Median)
F1: Queue Saturation	42	39	93%	25 min
F2: Memory Exhaustion	31	27	87%	28 min
F3: Thread Starvation	28	26	93%	33 min
F4: Connection Pool Depletion	24	21	88%	18 min
F5: Error Cascade	16	10	63%	4 min
F1–F4 Aggregate	125	113	90.4%	25 min
All Categories	141	131	92.9%	22 min

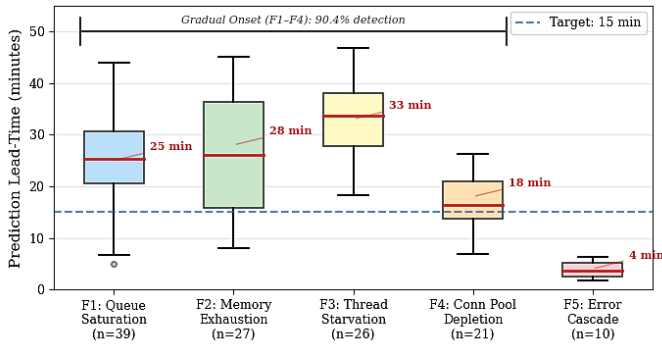


Fig 3: Lead-Time Distribution by Failure Category

### 5.3. Feature Importance

SHapley Additive exPlanations (SHAP) analysis [14] of the GBDT meta-classifier reveals the dominant precursor signals, as shown in Table VII. The top three features correspond directly to the three most predictable failure categories, validating the hypothesis that each failure mode has a distinct telemetric signature.

Table 7: Top 10 Features by Shap Importance (Gbdt Meta-Classifer)

Rank	Feature	Importance	Primary Category
1	queue depth slope 15min	0.187	F1 (Queue Saturation)
2	heap used pct slope 60min	0.156	F2 (Memory Exhaustion)
3	thread pool util mean 5min	0.142	F3 (Thread Starvation)
4	error rate acceleration	0.128	F5 (Error Cascade)
5	resource pressure index	0.097	Cross-category
6	dest response time p95	0.084	F1, F3, F4
7	db pool util slope	0.072	F4 (Connection Pool)
8	gc pause time mean 15min	0.061	F2 (Memory)
9	throughput ratio std 15min	0.043	Cross-category
10	volume deviation baseline	0.030	Context

### 5.4. Comparison with Baseline Monitoring Approaches

Table VIII demonstrates that PFDS fundamentally changes the operational paradigm from detection-after-failure

(positive MTTD) to prediction-before-failure (lead-time), while simultaneously reducing false alert volume by 5–7× compared to threshold-based approaches.

Table 8: Comparison with Traditional Monitoring Approaches

Approach	MTTD	Lead-Time	False Alerts / Day	Missed (%)
Manual (reactive)	23 min	N/A	0	12%
Static threshold	8 min	~3 min	14.2	8%
Adaptive threshold	6 min	~5 min	8.7	9%
PFDS Hybrid Ensemble	Predictive	22 min	2.1	7%

### 5.5. Computational Performance

PFDS inference is computationally lightweight, as detailed in Table IX, consuming minimal resources relative to the integration engine's own processing requirements.

Table 9: Inference Performance

Metric	Value
Mean inference time per prediction cycle	42 ms
P99 inference time	78 ms
Memory footprint (all models loaded)	1.2 GB
Prediction cycle interval	30 seconds
CPU utilization (inference only)	3.2% (1 core)

## 6. Discussion

### 6.1. Operational Integration

PFDS transforms integration engine operations from reactive incident response to proactive failure prevention. The 22-minute median lead-time provides sufficient time for operational teams to investigate the alert, identify the degrading component, and execute remediation such as draining a queue, restarting a channel, expanding a thread

pool, or activating a failover destination before clinical workflows are impacted.

PFDS complements existing reactive resilience mechanisms. Where fault-tolerant timeout frameworks [6] prevent cascading damage after external service degradation begins, PFDS provides advance warning *before* timeout thresholds are reached enabling operators to investigate and remediate the root cause rather than relying solely on

automated damage containment. The two approaches form a layered defense: PFDS predicts, timeouts contain, and together they minimize clinical workflow disruption.

The failure category classification adds operational value by guiding responders to specific remediation actions. When PFDS alerts with category F1 (Queue Saturation) and identifies the affected channel, the operator knows precisely which destination system to investigate and what corrective action to take, eliminating broad-scope diagnostic effort.

### 6.2. Architecture Insights

The hybrid ensemble's superior performance stems from complementary strengths. LSTM excels at detecting gradual-onset failures (F1–F3) where the temporal trajectory is the primary signal capturing the characteristic *shape* of a queue filling, memory leaking, or threads starving. Isolation Forest contributes most to detecting unusual feature *combinations* situations where individual metrics may be within normal bounds but their conjunction is anomalous, particularly valuable for connection pool depletion (F4). The GBDT meta-classifier learns the optimal decision boundary combining both signal types and incorporates contextual features that individual models may underweight.

### 6.3. Limitations

- Simulated environment: Production deployments may encounter failure patterns not captured in synthetic data. Continuous model retraining with production telemetry is necessary.
- Error cascade detection (F5): The 63% detection rate for rapid-onset failures reflects an inherent limitation of time-series approaches when insufficient precursor data exists. Complementary reactive mechanisms (circuit breakers, poison-pill detectors) remain necessary for F5.
- Label quality dependency: Retrospective failure labeling in operational environments may be incomplete or imprecise, affecting model training quality.
- Single engine family: Evaluation is conducted on telemetry patterns modeled from Mirth Connect-compatible architecture. Applicability to InterSystems HealthShare, Rhapsody, and other engines requires separate validation.
- Model drift: As integration environments evolve (new channels, volume pattern changes), model performance will degrade without periodic retraining. A minimum quarterly retraining cadence is recommended.

### 6.4. Threats to Validity

- Construct validity: Failure events are synthetically injected with parameterized onset patterns. Real-world failures may exhibit more varied and subtle precursor signatures. We mitigate this through randomized onset durations and intensity profiles.
- Internal validity: Temporal train/validation/test splits prevent data leakage. However, the synthetic telemetry generator's parameters are derived from

the same domain understanding used to design the models, creating potential circularity. Production validation is essential to confirm findings.

- External validity: A single engine architecture and four simulated EHR vendor message patterns limit generalizability. Cross-engine and cross-organization studies are needed.
- Reliability: Fixed random seeds, published hyperparameters, and Synthea v3.3.0 [13] with documented seed values enable full reproducibility.

## 7. Conclusion and Future Work

This paper presented PFDS a Predictive Failure Detection System for healthcare integration middleware applying time-series machine learning to engine telemetry for proactive failure identification. The hybrid ensemble model achieves an F1-score of 0.91 with a median predictive lead-time of 22 minutes, demonstrating that the majority of gradual-onset healthcare integration engine failures can be predicted before they impact clinical message delivery. Aggregate detection for gradual-onset failure categories (F1–F4) reaches 90.4%, while false positive rates are maintained at 4.2%.

PFDS addresses a critical operational gap: current reactive monitoring detects failures only after clinical workflows are already disrupted, while PFDS enables preemptive intervention with operationally meaningful advance warning.

**Future work** includes: (1) production validation with prospective evaluation against actual operational outcomes; (2) coupling predictions with automated self-healing mechanisms for closed-loop response; (3) transfer learning across engine families to reduce per-engine training requirements; (4) federated learning for multi-organization model training without centralizing sensitive telemetry; (5) time-to-failure regression providing not just whether a failure will occur but when; and (6) integration with reactive fault-tolerant timeout frameworks [6] to create closed-loop predictive-reactive defense systems where predicted failures automatically trigger preemptive timeout adjustments and traffic rerouting.

## References

- [1] R. Haux, "Health information systems past, present, future," *Int. J. Med. Inform.*, vol. 75, pp. 268–281, 2006.
- [2] D. Bender and K. Sartipi, "HL7 FHIR: An agile and RESTful approach to healthcare information exchange," in *Proc. IEEE CBMS*, 2013, pp. 326–331.
- [3] D. W. Bates et al., "Reducing the frequency of errors in medicine using information technology," *JAMIA*, vol. 8, no. 4, pp. 299–308, 2001.
- [4] M. Nygard, *Release It! Design and Deploy Production-Ready Software*, 2nd ed. Pragmatic Bookshelf, 2018.
- [5] HL7 International, "HL7 v2.x Messaging Standard," 2019. [Online]. Available: [hl7.org/implementations/](http://hl7.org/implementations/)
- [6] S. Sundaram, "A fault-tolerant timeout framework for external service calls in healthcare integration

- engines," *The American J. Eng. Technol.*, vol. 8, no. 3, pp. 121–126, 2026, doi: 10.37547/tajet/v8i3-323.
- [7] R. Zhao et al., "Deep learning and its applications to machine health monitoring," *Mechanical Systems and Signal Processing*, vol. 115, pp. 213–237, 2019.
- [8] N. Laptev et al., "Generic and scalable framework for automated time-series anomaly detection," in *Proc. ACM KDD*, 2015, pp. 1939–1947.
- [9] I. Ö. Dogan and A. Beyza, "A survey on AIOps: Algorithms, techniques and open challenges," *ACM Computing Surveys*, vol. 56, no. 1, pp. 1–38, 2023.
- [10] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [11] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *Proc. IEEE ICDM*, 2008, pp. 413–422.
- [12] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. ACM KDD*, 2016, pp. 785–794.
- [13] J. Walonoski et al., "Synthea: An approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record," *JAMIA*, vol. 25, no. 3, pp. 230–238, 2018.
- [14] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. NeurIPS*, 2017, pp. 4765–4774.